White Paper

Maximizing Data Center Efficiency with F5 BIG-IP AAM

Corporate networks are overburdened by the number of applications they serve and the increasing bandwidth demands of the data being transferred. The move from text to video content continues, for example, and the performance implications of these changes are large. Organizations can and should take advantage of intelligent bandwidth management and processing offload to keep network performance in top shape.

by Don MacVittie
Technical Marketing Manager II

Contents

Introduction

Network usage has been increasing for a very long time now. Be it over an internal network or external, VPN or public Internet, more people are trying to access more applications and data, and those applications and data consume greater bandwidth. From increasing web page sizes to streaming high-definition video, the curve of network bandwidth usage just continues to go up.

Consider that the average home page has exceeded 1 MB in size, according to HTTP Archive. Video streaming has surpassed peer-to-peer in terms of bandwidth usage, according to *Ars Technica* (see the sidebar), and worldwide Internet users grew by 566 percent from 2000 to 2012, according to InternetWorldStats.com. That is a lot of growth in three different dimensions.

The cost of this trend is often disassociated with the processes that cause the growth, though. Acquiring enough bandwidth and CPU cycles to serve the needs of the business tends to be addressed as a cost of doing business, as opposed to a variable that can be controlled. There are a variety of reasons for this, but as the given cost of doing business increases year after year, eventually the infrastructure overhead must be addressed.

While there will always be network and server overhead if IT is to function, the size of that overhead can be controlled and even significantly reduced by the use of intelligent architecture.

F5® BIG-IP® Application Acceleration Manager™ (AAM) is one such intelligent architecture. Able to reduce the bandwidth required to serve web-based applications, improve video streaming, reduce the associate bandwidth costs, and offload CPU intensive operations from servers in the data center (DC), BIG-IP AAM can control the growth of bandwidth needs and reduce server loads, delaying costly upgrades of both infrastructure components. While doing so, it can also optimize client-side connections (web, VPN), and back-end connections (DC to DC or DC to cloud) to improve performance of all network-based applications and delivery of all network-based resources.

> "Lots of applications are driving this growth, but most notable is video. Video surpassed peer-to-peer file sharing as the largest type of Internet traffic in 2010. It's expected to account for more than 50 percent of consumer Internet traffic by sometime this year. By 2015, on-demand video traffic will be the equivalent of three billion DVDs per month, and one million minutes worth of video will cross global IP networks every second."
>
> —John Brodkin, "Broadband Explosion: As Internet Use Soars, Can Bottlenecks Be Averted?" in *Ars Technica*, May 1, 2012

Bandwidth Growth and Change

There is no doubt that bandwidth usage—both at the global and corporate levels—is growing. As the World Wide Web becomes a more rich experience, standard web pages are growing in size, while adoption of video as a primary message delivery mechanism is growing as well, with payloads that are many times larger than the average web page. *Ars Technica* estimated in May 2012 that video would consume 50 percent of Internet traffic that year. Just for clarity, that 50 percent usage is in bandwidth, which means there are still far fewer videos being watched than other traffic on the Internet because video streams tend to be much larger and individually consume more bandwidth. Add in the growing tendency toward high definition (HD) streaming of videos, and the amount of network bandwidth that video will consume over time is on an exponential growth pattern.

And that means a lot of things, but mostly it means increasing bandwidth requirements. This growth is likely to be mirrored inside the organization as training and other corporate internal items are moved increasingly to high definition video.

Recently, Apple introduced HTTP live streaming (HLS), which has taken off across the mobile space and will likely become the pervasive delivery mechanism for all video. HLS breaks a video file into short segments and uses a playlist to deliver the segments. This introduces the ability to adapt more fluidly to the changing network environments that mobile delivery entails, offering a more agile response to bandwidth changes. HLS allows the client to dynamically choose which playlist the next segment should come from, meaning a fully enabled client can switch from high definition to low bandwidth settings between segments, keeping streaming going even if network conditions degrade. HLS is also designed to be run from a normal web server without specialized software on the server side. This makes delivery of video more attainable for shops on a budget and more maintainable for those that need video but for whatever reason do not want to pay a full-time employee to manage specialized video servers.

Since mobile delivery can entail higher latency and a higher error rate, retransmissions and renegotiations can be a problem compounded when there are many of them over the same network. On the data center end, this problem means lost bandwidth spent just to keep the client connected, and on the client end, it delays delivery of desired content.

> "HTTP Live Streaming lets you send audio and video over HTTP from an ordinary web server for playback on iOS-based devices—including iPhone, iPad, iPod touch, and Apple TV—and on desktop computers (OS X). HTTP Live Streaming supports both live broadcasts and prerecorded content (video on demand). HTTP Live Streaming supports multiple alternate streams at different bit rates, and the client software can switch streams intelligently as network bandwidth changes. HTTP Live Streaming also provides for media encryption and user authentication over HTTPS, allowing publishers to protect their work."
>
> —HTTP Live Streaming Overview, Apple Developer Library

Forward error correction (FEC) sends error correction information along with the data to give the client the ability to fix errors introduced on noisy lines.

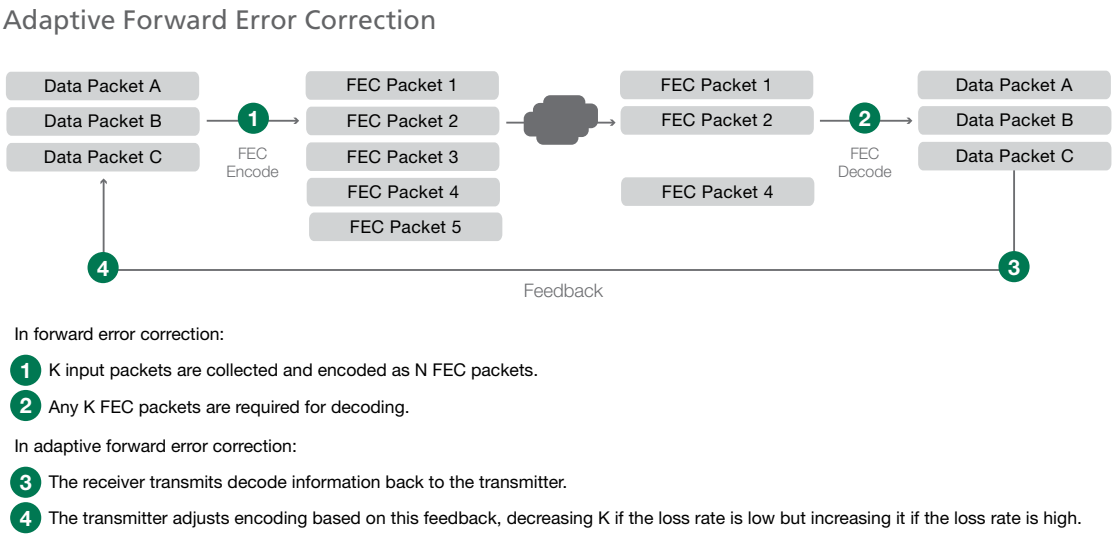

Figure 1: FEC in network communications

The benefits are a reduced number of round trips to reconstruct data and a more reliable communications stream. When latency is high and noise is low, the savings can be immense.

At the same time as bandwidth requirements have grown and the quality of networks has dipped (due to an increase in mobile traffic), SSL is becoming more prevalent for sites of all kinds, and many organizations are starting to use an always-on encryption methodology to protect their sensitive data from backdoor attacks from non-sensitive sites. This increases the load on data center infrastructure and limits visibility into traffic that is passing into and out of the network.

Data that's sent frequently, like the corporate logo or the first few HLS segments of a popular video, is a burden on the servers simply because they're serving the exact same thing over and over, rather than specializing in delivering the more valuable, calculated parts of a web page. The number of those parts is growing, too, as web sites are standardized and the volume of video continues to grow.

The BIG-IP AAM Solution

The best solution to these issues is one that is designed from the ground up for high-performance networking—one that can intelligently handle the growing array of protocols and adapt application delivery to the prevailing connection conditions. And one that can address networked application performance issues from several different approaches.

The F5 solution to this problem is BIG-IP AAM, a one-stop approach to application acceleration needs. Whether customer connections are taking too much bandwidth or back-end services like replication are causing network congestion, BIG-IP AAM can help. With client-side acceleration that includes support for HLS and FEC, and back-end synchronous support for replication and other DC-to-DC communications, BIG-IP AAM has the answers to improve application performance.

BIG-IP AAM also offloads several CPU-intensive operations from data center servers. Whether the work is encryption or compression, AAM can do it and leave servers to do what they do best—serve up applications. The savings in terms of delayed upgrades of server hardware can be impressive.

With support for SPDY, connection pooling to servers to reduce TCP overhead, various outbound TCP optimizations, and DC-to-DC tunneling with both compression and deduplication, BIG-IP AAM optimizes all facets of network communication, improving responsiveness and delaying the purchase of increased bandwidth.

In the case of HLS, BIG-IP AAM caches content frequently requested, and its advanced cache algorithms keep cache evictions down. Combining this functionality with the other acceleration features that touch all levels of the optimization problem—from TCP connection retries through HTTP compression—offers a complete solution to tough application delivery problems.

Conclusion

Network traffic is continuing to grow, both in volume of requests and the size of responses. Organizations need a comprehensive solution to the entire array of challenges this trend creates. While implementing one form of application acceleration may solve some of the issues a modern application faces, a comprehensive set of optimizations that can be directed at each of the major pain points is necessary for a long-term solution.

F5 BIG-IP AAM offers that comprehensive solution, addressing enhancements that range from TCP optimizations to video streaming and replication issues. By offloading burdensome tasks from data center servers, it also improves performance at the source of the response, with benefits in responses time and hardware upgrade frequency.

From video delivery to front end optimization, from WAN optimization to mobile content optimization, BIG-IP AAM comprehensively addresses application delivery performance needs and helps delay costly hardware upgrades or reduce expensive cloud throughput growth.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

©2013 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. WP-3950 0613